

Active Learning of Multiple Source Multiple Destination Topologies

Pegah Sattari, Maciej Kurant, Animashree Anandkumar, Athina Markopoulou

Department of EECS, University of California, Irvine

{psattari, mkurant, a.anandkumar, athina}@uci.edu

Michael Rabbat

Dept. of ECE, McGill University

michael.rabbat@mcgill.ca

Abstract—We consider the problem of inferring the topology of an M -by- N network by sending probes between M sources and N receivers. Prior work has shown that this problem can be decomposed into two parts: first, infer smaller subnetwork components (*i.e.*, 1-by- N 's or 2-by-2's) and then merge these components to identify the M -by- N topology. In this paper, we focus on the second part. In particular, we assume that a 1-by- N topology is given and that all 2-by-2 components can be queried and learned using end-to-end probes. The problem is *which* 2-by-2's to query and *how* to merge them with the 1-by- N , so as to exactly identify the 2-by- N topology, and optimize a number of performance metrics including measurement traffic, time complexity, and memory usage. We provide a lower bound, $\lceil \frac{N}{2} \rceil$, on the number of 2-by-2's required by any active learning algorithm and propose two greedy algorithms. The first algorithm follows a bottom-up approach: at every step, it selects two receivers, queries the corresponding 2-by-2, and merges it with the given 1-by- N ; it requires exactly $N - 1$ steps, which is much less than all $\binom{N}{2}$ possible 2-by-2's. The second algorithm follows the framework of multiple hypothesis testing, in particular Generalized Binary Search (GBS). Simulation results over synthetic and realistic topologies demonstrate that both algorithms correctly identify the 2-by- N topology and are near-optimal, but the bottom-up approach is more efficient in practice.

I. INTRODUCTION

Knowledge of network topology is important for network management, diagnosis, operation, security, and performance optimization [1–6]. In this paper, we consider a tomographic approach to topology inference, which assumes no cooperation from intermediate nodes and relies on end-to-end probes to infer internal network characteristics, including topology [4]. Typically, multicast or unicast probes are sent/received between sets of sources/receivers at the edge of the network, and the topology is inferred based on the number and order of received probes, or more generally, using some metric or correlation structure. An important performance metric is measurement bandwidth overhead: it is desirable to accurately infer the topology using a small number of probes.

In this paper, we focus on the problem of multiple-source multiple-destination topology inference: our goal is to infer the internal network (M -by- N) topology by sending probes between M sources and N receivers at the edge of the network. Prior work [1–3] has shown that this problem can be decomposed into two parts: first, infer smaller subnetwork components (*e.g.*, multiple 1-by- N 's or 2-by-2's) and then merge them to identify the entire M -by- N topology.

Significant progress has been made over the past years on the decomposition and the first part of the problem, *i.e.*, inferring smaller components (1-by- N 's or 2-by-2's) using active

probes. One body of work developed techniques for inferring 1-by- N (*i.e.*, single-source tree) topologies using end-to-end measurements [7–15]. Follow-up work [1–3] showed that an M -by- N topology can be decomposed into/reconstructed from a number of two-source, two-receiver (2-by-2) subnetwork components or “quartets”. In [1, 2], a practical scheme was proposed to distinguish between some quartet topologies using back-to-back unicast probes. In our recent work [16, 17], we proposed a method to exactly identify the topology of a quartet in networks with multicast and network coding capabilities.

In this paper, we focus on the second part of the problem, namely selecting and merging smaller subnetwork components to exactly identify the M -by- N , which has received significantly less attention than the first part. Existing approaches developed for merging the quartets [1, 3] have several limitations, including not being able to exactly identify the M -by- N topology and/or being inefficient (*e.g.*, requiring to send probes over all $\binom{N}{2}$ possible quartets). In this paper, we formulate the problem as active learning, characterize its complexity, and follow principled approaches for designing efficient algorithms to solve it. This complexity is important from both theoretical (a fundamental property of the topology inference problem) and practical (it determines the measurement bandwidth overhead, running time and memory usage) points of view. These costs can become particularly important when we need to infer large or dynamic topologies using active measurements.

More specifically, we start from the problem of 2-by- N topology inference, which is an important special case and can then be used as building block for inferring an M -by- N . Consistently with [1], we assume that a (static) 1-by- N topology is known (*e.g.*, using one of the methods in [4, 7–15, 18]) and that the topology of a quartet component can be queried and learned, if so desired (*e.g.*, using end-to-end probes and some of the methods in [1, 2, 16, 17, 19–23]¹). The problem then becomes one of active learning: *which* quartets to query and *how* to merge them with the given 1-by- N , so as to exactly identify the 2-by- N and optimize several performance metrics including measurement bandwidth, merging complexity and memory usage. Our contributions are as follows:

1) We provide a lower bound of $\lceil \frac{N}{2} \rceil$ on the number of quartets required by *any* active learning algorithm in order to identify the 2-by- N . This characterizes the inherent complexity of the problem and also serves as a rough baseline for assessing

¹Other techniques may also be developed in the future: this is still an active research area. But this is out of the scope of this paper (see Section III).

the performance of practical algorithms.

2) We design an efficient merging algorithm that follows a greedy bottom-up approach and provably identifies the 2-by- N by querying exactly $N - 1$ quartets. From the active probing perspective, this is attractive since only $N - 1$ queries are required, which is much lower than all $\binom{N}{2}$ possible queries.

3) We also formulate the problem within the framework of multiple hypothesis testing and develop an active learning algorithm based on Generalized Binary Search (GBS).

We compare the two algorithms to each other and to the lower bound via simulations over synthetic and realistic topologies. The results show that both algorithms can exactly identify the topology and are near-optimal in terms of active measurement bandwidth. Between the two, the bottom-up algorithm is very efficient in terms of running time and memory usage, and thus recommended for practical implementation.

The rest of the paper is organized as follows. Section II summarizes related work. Section III provides the problem statement and terminology. Section IV provides a lower bound on the number of quartets required by any algorithm. Section V proposes an efficient bottom-up algorithm and analyzes its correctness and performance. Section VI proposes another greedy algorithm based on the GBS framework. Section VII evaluates the two algorithms through simulations. Section VIII discusses possible extensions. Section IX concludes the paper.

II. RELATED WORK

There is a large body of prior work on inference of network topology. The most closely related to this paper are the ones using active measurements and network tomography.

One family of techniques relies on cooperation of nodes in the middle of the network, and uses *traceroute* [20–23] measurements to collect the ids of nodes along paths. However, some nodes may not respond and nodes often have multiple network interfaces (ids). Thus, *traceroute*-based methods must deal with missing/incomplete data and alias problems.

Unlike *traceroute*, tomographic approaches do not rely on responses from intermediate nodes, but only on end-to-end measurements. A survey of *network tomography* can be found in [4]. Most tomographic approaches rely on probes sent from a single source in a tree topology [7–15] and feed the number, order, or a monotonic property of received probes as input to statistical signal-processing techniques.

In [1–3], the authors formulated the multiple source multiple destination (M -by- N) tomography problem by sending probes between M sources and N receivers. It was shown that an M -by- N network can be decomposed into a collection of 2-by-2 components, also referred to as quartets [5, 6]. Coordinated transmission of back-to-back unicast probes from 2 sources and packet arrival order measurements at the 2 receivers were used to infer some information about the quartet topology. Assuming knowledge of M 1-by- N topologies and the quartets, it was also shown how to merge a second source's 1-by- N tree with the first one. The resulting M -by- N is not exact, but bounds were provided on the locations of the points where the two 1-by- N trees merge with each other. This approach also requires a large number of probes for statistical significance, similar to many other methods [7–11]. Compared to [1], our work is different in that (i) we assume perfect knowledge of the




Fig. 1. An example 2-by-4 topology. Solid lines and branching points $B_{i,j}$'s depict $G_{S_1 \times \mathcal{R}}$. J_i is a joining point, where P_{2i} (shown by dashed lines) joins $G_{S_1 \times \mathcal{R}}$. An example quartet is the part of the network connecting S_1, S_2 to R_1, R_2 , which is type 1 since both J_1, J_2 lie above the branching point of R_1, R_2 in $G_{S_1 \times \mathcal{R}}$, i.e., $B_{1,2}$.

quartets, thus we identify the topology accurately; (ii) we focus on the efficiency of active learning, i.e., selecting and merging the quartets, which has not been studied before. To the best of our knowledge, the only other merging algorithm proposed in the literature is [1, 3]. However, the merging was not efficient since all possible quartets were queried exhaustively.

In our prior work [16, 17], we revisited the problem of topology inference using end-to-end probes in networks where internal nodes are equipped with multicast and network coding capabilities. We built on [1] and extended it, using network coding at internal nodes to deterministically distinguish among all possible quartet topologies, which was not possible before. While in [16, 17], we focused on inferring quartets fast and accurately, here we assume that any quartet can be queried and learned, and focus on efficiently selecting and merging quartets to infer the larger topology. To the best of our knowledge, we are the first to look at this aspect of the problem.

There also exists a rich body of work on *multiple hypothesis testing*. One of the contributions of this paper is to formulate this problem in that framework and design an algorithm based on GBS [24–26], which we describe in detail in Section VI.

Topology inference problems have also been studied in the context of *phylogenetic trees* [27, 28]. [6] built on [28] and proposed robust algorithms for multiple source tree topology inference. [5] inferred the topology of sparse random graphs. However, the quartet structures and the way we measure them are different in our case due to the nature of active probing in network tomography (see problem formulation in Section III).

III. PROBLEM STATEMENT

M -by- N Topology to be inferred. Consider an M -by- N topology as a directed acyclic graph (DAG), between M source nodes $\mathcal{S} = \{S_1, \dots, S_M\}$ and N receivers $\mathcal{R} = \{R_1, \dots, R_N\}$. We denote this M -by- N topology by $G_{\mathcal{S} \times \mathcal{R}}$. Note that $G_{S_i \times \mathcal{R}}$, $i = 1, \dots, M$, is a 1-by- N tree. Similar to [1–3], we assume that a predetermined routing policy maps each source-destination pair to a unique route from the source to the destination. This implies the following three properties, first stated in [1]:²

²These assumptions are realistic, the same as in [1–3], and consistent with the destination-based routing used in the Internet: each router decides the next hop taken by a packet using a routing table lookup on the destination address. We further assume that the network does not employ load balancing.

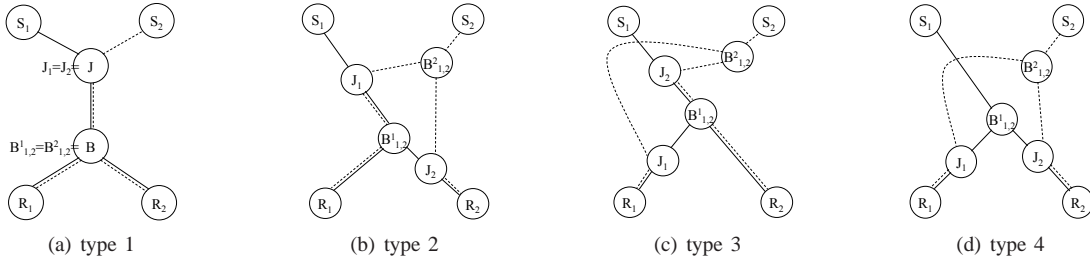


Fig. 2. The four possible types of a quartet (2-by-2 subnetwork component). There are two sources S_1, S_2 multicasting packets x_1, x_2 to two receivers R_1, R_2 . All links are directed downwards, but arrowheads are omitted to avoid cluttering. (The 1-by-2 topology of S_1 is a tree composed of $S_1, B_{1,2}^1, R_1, R_2$. Similarly, the 1-by-2 tree rooted at S_2 is $S_2, B_{1,2}^2, R_1, R_2$. J_1 and J_2 are joining points, where paths from S_2 to R_1 and R_2 join/merge with S_1 's tree.)

- A1 For every source S_i and every receiver R_j , there is a unique path P_{ij} .
- A2 Two paths P_{ij} and P_{ik} , $j \neq k$, branch at a *branching point* B , and they never merge again.
- A3 Two paths P_{ik} and P_{jk} , $i \neq j$, merge at a *joining point* J , and they never split again.

We are interested in inferring the logical topology³, defined by the branching and joining points defined above. We present most of our discussion in terms of $M = 2$, *i.e.*, inferring a 2-by- N topology $G_{S \times \mathcal{R}}$, $\mathcal{S} = \{S_1, S_2\}$; an M -by- N topology, $\mathcal{S} = \{S_1, \dots, S_M\}$, can then be constructed by merging smaller structures, as we describe in Section VIII.

Example 1: Fig. 1 illustrates an example 2-by- N topology with $N = 4$. The logical tree topology of S_1 is shown by solid lines and branching points $B_{i,j}$'s. Each J_i depicts a joining point, where the path from S_2 to receiver R_i (indicated by the dashed lines) joins the S_1 tree. For example, the path from S_2 to R_1 joins the S_1 tree at a point between $B_{1,3}$ and $B_{1,2}$, whereas the path to R_4 joins at a point above $B_{1,4}$. ■

Quartet Components. In [1], it has been shown that an M -by- N topology can be decomposed into a collection of 2-by-2 subnetwork components, which, in this paper, we call *quartets*, following the terminology in [5, 6]. Each quartet can be of four possible types, as shown in Fig. 2. We refer to Fig. 2 (a), (b), (c), and (d) as types 1, 2, 3, and 4, respectively.

In order to infer the type of a quartet between two sources S_1, S_2 and two receivers R_i, R_j , a set of probes must be sent from S_1, S_2 to R_i, R_j . The received probes can then be processed using techniques such as the ones developed in: [1, 2] (which distinguish type 1 from types 2, 3, 4 by sending back-to-back unicast probes); [16, 17] (which distinguish among all four types exploiting multicast and network coding); [19] (which can exactly infer the topology of a super-source to two receivers using network coding); traceroute [20–23] from the two sources to the two receivers; or other techniques that may be developed in the future, since this is still an active research area. We consider the design of these techniques to be out of the scope of this paper and we focus on their use by active learning algorithms to perform a *query*, *i.e.*, learn a quartet type by sending and processing a set of active probes.

Being able to query the type of a quartet enables inference of an M -by- N topology in two steps, as follows: first infer the

type of each quartet, and then merge these quartets to identify the original topology. Indeed, knowing the type of the quartet, we can use Fig. 2 to infer the relative location of joining and branching points. For example, knowing that the quartet is of type 1 implies that (i) the two joining points coincide $J_1 \equiv J_2$, (ii) the two branching points coincide $B_{1,2}^1 \equiv B_{1,2}^2$, and (iii) the joining point is above the branching point. Similar inferences can be made from the other types.

Problem Statement. Consistently with [1], we assume that $G_{S_1 \times \mathcal{R}}$ (*i.e.*, the 1-by- N tree topology rooted at S_1 , which contains only branching points) is known (*e.g.*, using one of the methods in [4, 7–15, 18]). We also assume that the type of the quartet between S_1 , a new source S_2 , and any two receivers can be queried and learned, as explained above.

Given (i) $G_{S_1 \times \mathcal{R}}$ and (ii) the ability to query the quartet type between S_1, S_2 , and any two receivers R_i, R_j , our goal is to *identify* all joining points, $\mathcal{J}_N = \{J_1, J_2, \dots, J_N\}$, where the paths from S_2 to each receiver join the tree describing paths from S_1 to the same set of receivers. Identifying a joining point J_i (for receiver R_i) means locating J_i on a single logical link, between two branching points on $G_{S_1 \times \mathcal{R}}$. *E.g.*, in Fig. 1, the path from S_2 to R_1 joins the S_1 tree at a point between nodes $B_{1,3}$ and $B_{1,2}$; *i.e.*, J_1 is located on the link $(B_{1,3}, B_{1,2})$.

We achieve this goal via active learning: we start from the given, static, 1-by- N topology $G_{S_1 \times \mathcal{R}}$ and proceed by updating it in steps. In each step, we select which quartet to query (*i.e.*, which two receivers to send probes to, from sources S_1, S_2)⁴, and learn its type (after sending and processing the received probes, we have essentially queried and learned the type of that quartet). We then merge this quartet with the known topology so far. We continue until identifying the entire 2-by- N . The goal is to exactly identify the 2-by- N topology while minimizing the number of queries (*i.e.*, set of probes sent to measure the quartets). This metric is important because it directly translates into measurement bandwidth. Additional performance metrics that it is desirable to keep low include: merging complexity and memory usage.

IV. LOWER BOUND

First, we provide a lower bound on the number of quartets required by any active learning algorithm to infer the 2-by- N . It clearly depends on the topology we want to infer and serves as a baseline for the performance of the proposed algorithms.

³A logical topology is obtained from a physical topology by ignoring nodes with in-degree = out-degree = 1. Such nodes cannot be identified and network tomography always focuses on inferring logical topologies.

⁴Since we focus on $M = 2$, *i.e.*, only two sources S_1 and S_2 , we represent the quartets (S_1, S_2, R_i, R_j) only by the receivers (R_i, R_j) for brevity.




Fig. 3. Two example 2-by- N topologies with $N = 4$. In (a), $\frac{N}{2}$ quartets are sufficient to identify the joining points, i.e., (R_1, R_2) and (R_3, R_4) . In (b), more than $\frac{N}{2}$ quartets are required, e.g., (R_1, R_2) , (R_1, R_3) , and (R_1, R_4) .

Theorem 4.1: Given $G_{S_1 \times \mathcal{R}}$, the number of quartets required to be queried by any algorithm in order to identify all the joining points in $G_{S \times \mathcal{R}}$, $\mathcal{S} = \{S_1, S_2\}$, is at least $\lceil \frac{N}{2} \rceil$.

Before proving the theorem, let us discuss some examples that illustrate the intuition and that this bound is not tight.

Example 2: Fig. 3(a) shows a 2-by- N topology with $N = 4$, which requires querying exactly $\frac{N}{2} = 2$ quartets in order to uniquely identify all the joining points. This is because, in this particular topology, knowing the types of (R_1, R_2) and (R_3, R_4) is sufficient for identifying all four joining points. Indeed, (R_1, R_2) is of type 4, which, according to Fig. 2, means that both J_1 and J_2 lie below $B_{1,2}$; also (R_3, R_4) is type 4, which means that both J_3 and J_4 are below $B_{3,4}$. Thus, each joining point is identified on a single logical link. ■

Example 3: Fig. 3(b) shows an example where $\frac{N}{2} = 2$ quartets are not sufficient and 3 quartets are needed to identify all joining points. There exist $\binom{4}{2} = 6$ possible quartets in this topology, from which $\binom{6}{2} = 15$ pairs of quartets can be selected; one can check that none of the 15 possible pairs can uniquely identify all joining points. For example, let us consider (R_1, R_2) . Since it is of type 1, Fig. 2 indicates that $J_1 \equiv J_2$ and both of them lie above $B_{1,2}$. But there is more than a single link above $B_{1,2}$; thus we continue by considering (R_1, R_3) . It is again of type 1, which means that $J_1 \equiv J_3$ is located above $B_{1,3}$. Thus, we go one step further and consider (R_1, R_4) . Since this is also type 1, $J_1 \equiv J_4$ lies above $B_{1,4}$. At this step, we only have a single link between S_1 and $B_{1,4}$ and thus, $J_1 \equiv J_2 \equiv J_3 \equiv J_4$ are all identified (depicted as J in Fig. 3(b)). Although there are other choices of triplets of quartets, in this topology, at least 3 quartets are required. ■

From these examples, one can see that the lower bound of $\lceil \frac{N}{2} \rceil$ is not tight and it is not achievable in every topology. Theorem 4.1 follows from the following lemma.

Lemma 4.2: In order for an algorithm to identify all joining points for all the receivers, each receiver needs to appear in the set of quartets queried by the algorithm at least once.

Proof: Assume that there exists a receiver R_i that has not been queried in any of the quartets. We show that even with complete knowledge of all other joining points, there exist at least two possible and feasible locations for J_i , as follows.

Location 1: J_i lies on the last incoming link to R_i , i.e., on the link between the parent of R_i in the S_1 tree (which from now on, we denote by $\text{parent}(R_i)$), and R_i . For example in




Fig. 4. Deletion and contraction of edge e_4 in a graph.

Fig. 3(a) and Fig. 3(b), assume that $R_i = R_2$; then Location 1 would be the link $(B_{1,2}, R_2)$. This is allowed by the routing assumptions in Section III because (1) there is a unique path P_{2i} ; (2) P_{2i} never merges with P_{2j} , $j \neq i$; and (3) P_{2i} merges with P_{1i} at J_i , and they continue together until they reach R_i .

Location 2: Define J_i as follows. On path P_{1i} , start at $\text{parent}(R_i)$ and move up towards S_1 , until the first link that does not fully overlap with any P_{2j} , $j \neq i$. Place J_i on that link. For example in Fig. 3(a), Location 2 for J_2 would be the link $(B_{1,3}, B_{1,2})$; whereas in Fig. 3(b), it would be $(S_1, B_{1,4})$. This location is also allowed by the assumptions in Section III:

- A1 There is a unique path P_{2i} .
- A2 For every $j \neq i$, the two paths P_{2i} and P_{2j} never join after they branch. Indeed, if J_j is located above J_i on P_{1i} , then this is guaranteed by the construction of J_i . In contrast, J_j cannot be located below J_i on P_{1i} since this would imply the violation of A2 even before adding J_i .
- A3 P_{2i} merges with P_{1i} at J_i and they never split.

Thus, both Location 1 and Location 2 are valid for J_i , according to the routing assumptions, and J_i cannot be uniquely identified. Therefore, R_i needs to be queried at least once. ■

Theorem 4.1 follows from the following reasoning: each quartet involves two receivers, and thus, at least $\lceil \frac{N}{2} \rceil$ quartets are required for each receiver to appear in the set of quartets queried by the algorithm at least once.

V. A BOTTOM-UP GREEDY ALGORITHM

In this section, we design a greedy algorithm that given $G_{S_1 \times \mathcal{R}}$, and the ability to query the type of any quartet, it is able to identify all N joining points where $G_{S_2 \times \mathcal{R}}$ merges with $G_{S_1 \times \mathcal{R}}$, i.e., the entire 2-by- N topology, in $N - 1$ steps.

Let every edge e in $G_{S_1 \times \mathcal{R}}$ have a unique name: $\text{label}(e)$. In our algorithm, we use two operations “edge deletion” and “edge contraction”, depicted in Fig. 4 and defined as follows.

Definition 1: Deleting edge (u, v) , entails taking that edge out of the graph while the end-nodes u and v , and the labels of the remaining edges in the graph remain unchanged.

Definition 2: Contracting edge (u, v) into node w , consists of deleting that edge and merging u and v into a single node w . The labels of the remaining edges do not change (although nodes may be renamed to w).

The algorithm is described in Alg. 1. It starts from the S_1 tree ($G_{S_1 \times \mathcal{R}}$) and proceeds by selecting one quartet to

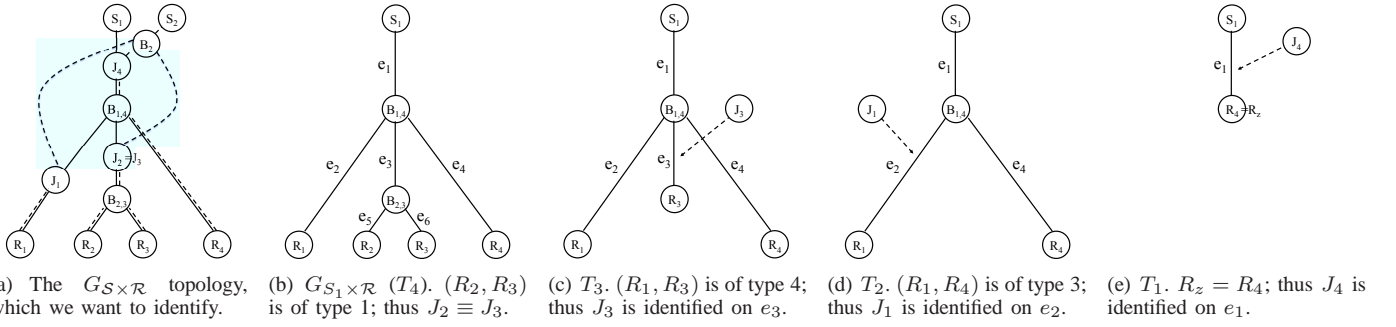


Fig. 5. The steps (b), (c), (d), and (e), performed by Alg. 1 to identify the 2-by- N topology in (a). The output of the algorithm is $J = [e_2, e_3, e_3, e_1]$.

Algorithm 1 Bottom-up merging algorithm: it starts from $G_{S_1 \times \mathcal{R}}$, selects the quartets sequentially, queries their types, and merges them until identifying all joining points \mathcal{J}_N .

```

1: Let  $J$  be a vector of length  $N$  of edge labels, which represents the
   locations of the joining points.
2: while  $|\mathcal{R}| > 1$  do
3:   Pick any two receivers  $R_i, R_j$  in  $G_{S_1 \times \mathcal{R}}$ , such that  $R_i$  and  $R_j$  are
   siblings; denote their parent by  $P$ .
4:   Query the type of  $(R_i, R_j)$ .
5:   switch  $(R_i, R_j)$  do
6:     case type 1 :
7:        $J_i \equiv J_j$ 
8:       Delete  $R_i$  and edge  $(P, R_i)$ .
9:       if  $\text{outdeg}(P)=1$  then
10:        Contract  $(P, R_j)$  into  $R_j$ .
11:     case type 2 :
12:        $J_j = \text{label}((P, R_j))$ 
13:       Delete  $R_j$  and edge  $(P, R_j)$ .
14:       if  $\text{outdeg}(P)=1$  then
15:        Contract  $(P, R_i)$  into  $R_i$ .
16:     case type 3 :
17:        $J_i = \text{label}((P, R_i))$ 
18:       Delete  $R_i$  and edge  $(P, R_i)$ .
19:       if  $\text{outdeg}(P)=1$  then
20:        Contract  $(P, R_j)$  into  $R_j$ .
21:     case type 4 :
22:        $J_j = \text{label}((P, R_j))$ 
23:       Delete  $R_j$  and edge  $(P, R_j)$ .
24:       if  $\text{outdeg}(P)=1$  then
25:        Contract  $(\text{parent}(P), P)$  into  $P$ .
26: /*There is one remaining receiver, which we call  $R_z$ .*/
27: Let  $J_z = \text{label}((\text{parent}(R_z), R_z))$ .
28: Output  $J$ .

```

query at each step (*i.e.*, 2 receivers R_i, R_j to send probes to, from sources S_1, S_2). The two receivers (R_i, R_j) in the selected quartet are sibling leaves. Based on the type of the selected quartet, Alg. 1 identifies exactly one joining point in one step. It then updates $G_{S_1 \times \mathcal{R}}$ by deleting the receiver whose joining point has been identified and the last incoming edge to that receiver. Furthermore, if a node of degree two appears in $G_{S_1 \times \mathcal{R}}$ as a result of this edge deletion, the algorithm eliminates that node by contracting the corresponding edge. The algorithm continues iteratively until there is one edge left, *i.e.*, all joining points are identified. This way, Alg. 1 identifies all joining points (where paths from S_2 to each receiver join the S_1 tree), one-by-one, proceeding from the bottom to the root of the tree. Next, we describe an illustrative example.

Example 4: Fig. 5(b)-(e) demonstrate the steps performed by Alg. 1 to identify the 2-by- N topology shown in Fig. 5(a).

The algorithm starts from $G_{S_1 \times \mathcal{R}}$ shown in Fig. 5(b); e_1, \dots, e_6 are the edge labels on this tree. The algorithm first selects (R_2, R_3) and queries its type. Since the answer is type 1, the algorithm assigns $J_2 \equiv J_3$, and deletes R_2 and e_5 . Since the degree of $B_{2,3}$ becomes 2, the algorithm contracts e_6 into R_3 .

In the second step shown in Fig. 5(c), Alg. 1 selects two sibling leaves (R_1, R_3) , randomly out of three possible pairs of siblings, and queries its type. Since it is type 4, the algorithm identifies J_3 on e_3 (which, together with the previous step, means that J_2 is also identified). It also deletes R_3 and e_3 . There is no contraction in this step as $B_{1,4}$'s degree is > 2 .

In the third step shown in Fig. 5(d), (R_1, R_4) is selected and queried; it is of type 3. Therefore, the algorithm identifies J_1 on e_2 , deletes R_1 and e_2 , and contracts e_4 into R_4 . Since there is only one receiver left, there are no more quartets to query; thus the algorithm exits the while loop and proceeds to the last step (line 26). For $R_z = R_4$, the algorithm identifies J_4 on e_1 , as shown in Fig. 5(e). The identified joining points agree with the real locations in $G_{S \times \mathcal{R}}$ topology in Fig. 5(a), which demonstrates the correctness of the algorithm. ■

A. Properties of Algorithm 1

Let $T_N = G_{S_1 \times \mathcal{R}}$ denote the logical tree from S_1 to all N receivers, which we assume to be known. In this section, we use the notation T_N to emphasize that this initial tree $G_{S_1 \times \mathcal{R}}$ contains N receivers. After each iteration through the while loop in Alg. 1, one receiver is deleted. We write T_k to denote the tree (rooted at S_1) obtained at the end of iteration $(N - k)$, at which point there are k receivers remaining. Let \mathcal{J}_k denote the set of joining points, which still remain to be identified after iteration $(N - k)$, *i.e.*, one for each remaining receiver.

Proposition 5.1: Let T_k and \mathcal{J}_k be given. The next iteration of Alg. 1 (lines 3–25) produces T_{k-1} and \mathcal{J}_{k-1} , which satisfy the following properties:

- 1) The S_1 topology is still a logical tree, and it has $k - 1$ receivers (*i.e.*, one receiver and its corresponding edge are deleted from T_k). Therefore, we denote it by T_{k-1} .
- 2) One joining point has been identified; therefore, the algorithm has $k - 1$ more joining points in \mathcal{J}_{k-1} to identify.
- 3) All joining points in \mathcal{J}_{k-1} are located on edges in T_{k-1} .

Proof: These properties follow directly from the operations performed by one step of Alg. 1:

- 1) In each iteration, a single receiver is eliminated from the tree. Consequently, the only node that can possibly have degree two (or out-degree one) after deleting the receiver is its parent,

P . However, after each deletion, Alg. 1 tests to see if P has out-degree 1, and if it does, then an additional contraction is performed so that the resulting tree, T_{k-1} , is still logical.

2) When (R_i, R_j) is of type 2, 3, or 4, we can see in lines 12, 17, and 22 of the algorithm, respectively, that one joining point is identified. When (R_i, R_j) is of type 1, line 7 assigns to R_i , the same joining point as R_j 's. Then, in line 8, R_i is deleted so that we do not create a loop by assigning J_i again to J_j later. Also, J_j eventually becomes identified, either in one of the other types (2, 3, or 4) in the while loop, or in the last line of the algorithm. Thus, we have \mathcal{J}_{k-1} after one step.

3) Alg. 1 changes T_k by 2 processes: edge deletion and edge contraction. We show that neither deletion nor contraction can eliminate an edge in T_k that contains a joining point in \mathcal{J}_{k-1} .

Deletion: Alg. 1 is constructed s.t. any edge deleted from the S_1 tree contains either no joining point (if (R_i, R_j) is type 1) or exactly one joining point, corresponding to the receiver being removed along with that edge (if (R_i, R_j) is type 2,3,4).

Contraction: An edge is contracted only when it does not contain any joining point, neither for R_i and R_j (see lines 9 – 10 for type 1, lines 14 – 15 for type 2, lines 19 – 20 for type 3, and lines 24 – 25 for type 4), nor for any other receivers (since (R_i, R_j) are sibling leaves, the contracted edge cannot contain any joining point for any other receiver.⁵) ■

The following theorem establishes the correctness and complexity of Algorithm 1.

Theorem 5.2: Alg. 1 terminates in N steps and correctly identifies all N joining points after querying $N - 1$ quartets.

Proof: The proof is via induction. In the beginning, $T_N = G_{S_1 \times \mathcal{R}}$ is a logical tree and according to Corollary 1 in [1], the joining points are identifiable using sufficient quartets. Our inductive step is one iteration of the while loop. First, note that there exist two sibling receivers at every step: it is enough to pick one of the lowest receivers (*i.e.*, a receiver with the largest distance from the source); it will always have a sibling because of the logical tree topology. The algorithm queries one quartet per step, identifies one joining point per step, and at the end of the step, it preserves properties 1, 2, and 3. The while loop terminates in $N - 1$ iterations and there is one additional step for R_z after the loop (which does not use any quartet). Thus, the algorithm terminates in N steps, and correctly identifies all N joining points by querying exactly $N - 1$ quartets. ■

Discussion. An important observation is that the $N - 1$ quartets are not known a priori, but are easily selected in a sequential way, as needed; this makes Alg. 1 easy to implement in practice using active probing. Another observation is about the running time: exactly $N - 1$ quartets need to be queried (by sending sets of probes). This is much less than the $\binom{N}{2}$ possible quartets queried by a brute-force approach [1, 3], but higher than the lower bound on the number of required quartets by any algorithm ($\lceil \frac{N}{2} \rceil$, Theorem 4.1). Therefore, Alg. 1 is not optimal, but it is simple, efficient, and provably correct.

VI. A GENERALIZED BINARY SEARCH ALGORITHM

A. Background on GBS

The GBS problem is defined as follows [24]. Consider a finite (potentially very large) collection of binary-valued

functions \mathcal{H} , called the “hypothesis space”, defined on a domain \mathcal{X} , called the “query space”. Each $h \in \mathcal{H}$ is a mapping from \mathcal{X} to $\{+1, -1\}$. Let $|\mathcal{H}|$ denote the cardinality of \mathcal{H} , *i.e.*, the total number of hypotheses. The functions $h \in \mathcal{H}$ are assumed to be unique, and one function, $h^* \in \mathcal{H}$, produces the correct binary labeling. h^* is assumed to be fixed but unknown. The goal is to determine h^* through as few queries from \mathcal{X} as possible. Thus, the queries need to be selected strategically in a sequential manner s.t. h^* is identified as quickly as possible.

This is an NP-complete problem [29]. A practical heuristic is given by a greedy algorithm called generalized binary search (GBS). In this section, we develop a GBS approach to our problem for the following reasons: (i) our problem is one of active learning and lends itself naturally to be posed in the GBS framework; (ii) GBS is a principled (although not optimal) approach with provable correctness and performance guarantees [24]; (iii) GBS can serve as a baseline for comparison with Alg. 1, in terms of the number of queries and complexity.

At each step, GBS selects a query that results in the most even split of the hypotheses under consideration into 2 subsets, responding $+1$ and -1 respectively, to the query. The correct response to the query eliminates one of these two subsets from further consideration. The work in [24] characterizes the worst-case number of queries required by GBS in order to identify the correct hypothesis h^* . The main result of [24] indicates that under certain conditions on the query and hypothesis spaces, the query complexity of GBS (*i.e.*, the minimum number of queries required by GBS to identify h^*) is near-optimal, *i.e.*, within a constant factor of $\log_2 |\mathcal{H}|$. The constant depends on two parameters c^* and k , defined in [24], and it is desirable that they are both as small as possible.

B. Merging Logical Topologies in the GBS Framework

In this section, we formulate our problem within the GBS framework. Consider a set of hypotheses \mathcal{H} , where each hypothesis $h \in \mathcal{H}$ is a configuration that results from placing each joining point J_i on an arbitrary link in the path P_{1i} in the S_1 tree. The query space \mathcal{X} is the set of all queries for all the quartets, where each query $x \in \mathcal{X}$ asks about the type of a quartet (R_i, R_j) . Since in our problem, each such query x has 4 possible answers (corresponding to the 4 quartet types), we need to modify our queries to make them consistent with the binary functions in the standard GBS framework. We assume that each query x consists of 4 subqueries, each of which asks whether (R_i, R_j) is of a specific type (1, 2, 3, or 4) or not; *i.e.*:

$$x = \begin{cases} \text{Is } (R_i, R_j) \text{ of type 1?} \\ \text{Is } (R_i, R_j) \text{ of type 2?} \\ \text{Is } (R_i, R_j) \text{ of type 3?} \\ \text{Is } (R_i, R_j) \text{ of type 4?} \end{cases}$$

The answer to each such subquery is binary, which is consistent with the GBS formulation. Of course, not all four subqueries are always required for a quartet; one would stop as soon as she gets the first “yes”, which would reveal the type of the quartet. Note, however, that we count the number of queries (not subqueries) as the performance metric of the GBS algorithm when comparing with Algorithm 1.

Our goal is to find the target hypothesis h^* , which is the configuration that results from the correct placement of the

⁵Alg. 1 selects *siblings* R_i, R_j at each step. Thus, there are only 2 potential candidates for the joining points that can be identified at this step: J_i, J_j .




Fig. 6. Four synthetic $G_{S_1} \times \mathcal{R}$ topologies used to compare the performance of Alg. 1 (the bottom-up approach) with Alg. 2 (the GBS approach).

Algorithm 2 GBS algorithm for identifying the joining points.

```

1: Let  $J = [0, 0, \dots, 0]$  be a vector of length  $N$ , which represents the
   locations of the joining points.
2: while  $\exists 0$  in  $J$  do
3:   Let  $wcB = []$  represent the worst case benefits for all the quartets.
4:   for each receiver  $R_i$  do
5:     for each receiver  $R_j, j > i$  do
6:       Let  $B_{i,j}$  be the lowest common ancestor of  $R_i, R_j$  in  $G_{S_1} \times \mathcal{R}$ 
7:       Let  $up_i \subset P_{1i}$  be the subset of  $P_{1i}$  located above  $B_{i,j}$ 
8:       Let  $up_j \subset P_{1j}$  be the subset of  $P_{1j}$  located above  $B_{i,j}$ 
9:       Let  $dn_i \subset P_{1i}$  be the subset of  $P_{1i}$  located below  $B_{i,j}$ 
10:      Let  $dn_j \subset P_{1j}$  be the subset of  $P_{1j}$  located below  $B_{i,j}$ 
11:       $type1\_B = \frac{|up_i|}{|P_{1i}|} \frac{|up_j|}{|P_{1j}|}$ 
12:       $type2\_B = \frac{|up_i|}{|P_{1i}|} \frac{|dn_j|}{|P_{1j}|}$ 
13:       $type3\_B = \frac{|dn_i|}{|P_{1i}|} \frac{|up_j|}{|P_{1j}|}$ 
14:       $type4\_B = \frac{|dn_i|}{|P_{1i}|} \frac{|dn_j|}{|P_{1j}|}$ 
15:       $wcB.append(\max\{type1\_B, type2\_B, type3\_B, type4\_B\})$ 
16:    $selectedQuartet = wcB.index(\min(wcB))$ 
17:   Let  $selectedQuartetType$  be the type of  $selectedQuartet$ .
18:   switch  $selectedQuartetType$  do
19:     case type 1 :
20:        $P_{1i} \leftarrow up_i$ 
21:        $P_{1j} \leftarrow up_j$ 
22:     case type 2 :
23:        $P_{1i} \leftarrow up_i$ 
24:        $P_{1j} \leftarrow dn_j$ 
25:     case type 3 :
26:        $P_{1i} \leftarrow dn_i$ 
27:        $P_{1j} \leftarrow up_j$ 
28:     case type 4 :
29:        $P_{1i} \leftarrow dn_i$ 
30:        $P_{1j} \leftarrow dn_j$ 
31:   if  $|P_{1i}| == 1$  then
32:      $J_i = P_{1i}$ 
33:   if  $|P_{1j}| == 1$  then
34:      $J_j = P_{1j}$ 
35: Output  $J$ .

```

joining points in the S_1 topology, using as few queries (*i.e.*, the knowledge of as few quartet types) as possible.⁶

Alg. 2 describes a greedy strategy based on GBS for determining h^* . In the beginning, there are $|\mathcal{H}|$ possible hypotheses. In each step, the algorithm selects the best (*i.e.*, maximally discriminating [24]) quartet to query as follows. By querying a quartet and learning its type, some information is obtained about the locations of two joining points. Thus, the number of feasible hypotheses, which agree with the constraints imposed

⁶More formally, h^* answers every query, for any pair of receivers, in accordance with the true 2-by- N topology. Mathematically, h^* is a mapping from queries to $\{+1, -1\}$, not a topology itself. However, there is a bijection between all 2-by- N logical topologies and corresponding mappings in \mathcal{H} , and therefore, knowing h^* is equivalent to knowing the 2-by- N topology.

by the quartets queried and learned so far, is reduced by a number, which depends on the topology in general. We call this number the *benefit* of the quartet. The best quartet to select to query is the one with maximum benefit. However, the benefit of each quartet becomes known only *after* it is queried. Thus, the algorithm considers all four possible types for every possible quartet, and focuses on the worst case benefit of that quartet, *i.e.*, the type that gives the minimum benefit. The best quartet to query is the one with maximum worst case benefit.

We denote the benefit of each type for a quartet (R_i, R_j) by $type1_B, \dots, type4_B$ in Alg. 2, and define it as follows. Each quartet type limits the number of candidate edges where J_i and J_j can be located on, in the way depicted in Fig. 2. The benefit of a type for (R_i, R_j) is the ratio of the number of edges where J_i and J_j can potentially be located on after learning this type, divided by the current number of candidate edges for the locations of J_i and J_j . The worst case (minimum) benefit of (R_i, R_j) results from the type for which this ratio is maximized, and the maximum of these worst case benefits over all quartets is given by the quartet with minimum ratio.

In order to provide an analytical upper bound on the number of quartets required by Alg. 2, one can try to use the main result of [24], which indicates that Alg. 2 requires $\log_2 |\mathcal{H}|$ quartets.⁷ However, we cannot compute $|\mathcal{H}|$ exactly in our problem; we can only provide a loose upper bound on that, which is $N!$.⁸ Thus, we get the bound of $\log N! \approx N \log N$ on the number of quartets required by Alg. 2, which is loose, and much larger than the $N - 1$ quartets of Alg. 1. The next section evaluates the performance of Algorithms 1, 2 via simulation.


VII. PERFORMANCE EVALUATION

A. Simulation Setup


We evaluate the two algorithms in simulations over both synthetic topologies (as shown in Fig. 6) and realistic topologies (as shown in Fig. 7). We compare them to each other as well as to the lower bound. The main performance metric of interest is the number of quartets queried in order to exactly infer the topology, which directly translates into *measurement overhead*. Additional metrics include the running time and the memory used by each algorithm.

⁷This is the best case, where the constants c^* and k in [24] are both as small as possible. In practice, there is an additional constant factor for $\log_2 |\mathcal{H}|$.

⁸The bound is obtained by starting from the S_1 tree and considering all possible placements of J_i on P_{1i} , $\forall i$. Fig. 6(c) shows that there are $N \times N \times (N - 1) \times \dots \times 2 \cong N!$ possible such placements. In practice, the routing assumptions in Section III impose some constraints on possible J_i locations. Also, the type of each quartet may rule out some types for the other quartets. Therefore, the exact $|\mathcal{H}|$ depends on the topology and we cannot compute it.



(a) A 2-by-16 topology from a LAN [3].



(b) A 2-by-16 topology from Exodus [23].

Fig. 7. Two realistic 2-by- N topologies used to compare the performance of Alg. 1 (Bottom-Up) with Alg. 2 (GBS). Solid lines indicate the paths taken by probes from S_1 ; dashed lines indicate the paths taken by probes from S_2 .

For the synthetic topologies, we illustrate only the 1-by- N tree topology of S_1 in Fig. 6. We consider the star topology, “perfect” and “tall” binary trees (referring to the topologies depicted in Fig. 6(b) and 6(c), respectively), and perfect ternary trees, for the $G_{S_1 \times \mathcal{R}}$ tree topology. Starting from this tree, we then create a 2-by- N topology, with sources S_1 and S_2 , by choosing the location of each joining point J_i (for receiver R_i) on a single logical link, selected uniformly at random, on P_{1i} in $G_{S_1 \times \mathcal{R}}$. For each $G_{S_1 \times \mathcal{R}}$ in Fig. 6, we consider 100 realizations of such random placements (resulting in different 2-by- N topologies) and report the average number of quartets required for these topologies in the next section.

For the realistic topologies, we show the complete 2-by- N topology in Fig. 7. Fig. 7(a) depicts a US University departmental LAN with 16 receivers, first used in [3]. Fig. 7(b) is a 2-by-16 directed acyclic graph (DAG), extracted from the Exodus topology, which is a large commercial ISP whose backbone map was inferred by the Rocketfuel project [23]. To generate this topology, we picked randomly two nodes of Exodus (nodes 5, 36) to be the sources, and selected all sixteen nodes to which both sources had routes to be the receivers. We then found the shortest path trees from each source to the receivers, and considered the overlap between these two trees.

Our experiments are conducted using Python implementations of Algorithms 1 and 2, which we have made available online [30]. They take as input any topology and return the number of quartets required by the two algorithms. Next, we summarize the simulation results.

B. Simulation Results (for the Number of Quartets)

When $G_{S_1 \times \mathcal{R}}$ is a star topology as depicted in Fig. 6(a), Alg. 2 always identifies the 2-by- N topology by querying only $\lceil \frac{N}{2} \rceil$ quartets, which is the lower bound; thus, it is optimal and performs better than Alg. 1, which requires $N - 1$ quartets.




Fig. 8. Simulation results for the average number of quartets required by Alg. 2 (GBS) to infer the 2-by- N when $G_{S_1 \times \mathcal{R}}$ is a perfect binary tree (Fig. 6(b)) of various sizes, $N = 4, \dots, 128$. The results are averaged over 100 realizations of random placements of the joining points. The standard deviation error bars (not shown) are comparable with the marker size.

When $G_{S_1 \times \mathcal{R}}$ is a perfect binary tree as shown in Fig. 6(b), Alg. 2 requires different numbers of quartets, between $\frac{N}{2}$ and N , in different 2-by- N topologies. However, as shown in Fig. 8, on average, Alg. 2 performs very close to Alg. 1, while being much more complex than Alg. 1.

Similar results are obtained for tall binary trees and perfect ternary trees. Due to lack of space, we omit the figures and report the results. When $G_{S_1 \times \mathcal{R}}$ is a tall binary tree as shown in Fig. 6(c), the number of quartets required by Alg. 2 varies depending on the quartet types in different 2-by- N topologies, but in our simulations on tall binary trees with $N > 100$ receivers, we observe that in at least 80% of the realizations, Alg. 2 requires the same number of quartets as Alg. 1. This percentage increases up to 99% in topologies with $N < 100$. When $G_{S_1 \times \mathcal{R}}$ is a perfect ternary tree, again on average, Alg. 2 performs close to Alg. 1, but for some topologies, Alg. 2 requires even more than N quartets.

For the realistic topologies in Fig. 7(a) and 7(b), Alg. 2 identifies both 2-by-16 topologies by querying 14 ($= N - 2$) quartets, while Alg. 1 requires $N - 1 = 15$ quartets.

Thus, in our simulations, we find that Alg. 2 only requires significantly fewer quartets than Alg. 1 for flat $G_{S_1 \times \mathcal{R}}$ topologies, such as the star in Fig. 6(a). In other topologies, such as binary/ternary trees or realistic topologies, Alg. 1 is generally preferred over Alg. 2, because it is simpler and identifies the joining points using the same number of quartets (or even fewer quartets in large topologies) as Alg. 2 (*i.e.*, $N - 1$).

C. Time and Space Complexity

1) *Time Complexity*: The time complexity of Alg. 2 ($O(N^3)$) is significantly higher than that of Alg. 1 ($O(N)$). The reason is that at each step, Alg. 1 only needs to select a pair of sibling receivers (any of them will do); while Alg. 2 calculates the worst case benefits of all the quartets, in order to pick the best one among them, which takes much longer.

As an example, for a single realization of our simulations when $G_{S_1 \times \mathcal{R}}$ is a perfect binary tree with 128 receivers, the running time of Alg. 2 is 19 seconds, while that of Alg. 1 is < 1 second. This is a big difference when we consider a large number of realizations as described in the previous section.

2) *Memory Usage*: The memory requirement of Alg. 2 is also much higher than that of Alg. 1. The reason is that Alg. 1 only requires to store the (modified version of the) graph at each step; while Alg. 2 requires to keep track of all the benefits and the worst case benefits for all the quartets, all the path updates for the location of each joining point, and so forth.

VIII. EXTENSIONS

Due to lack of space, we only briefly outline possible extensions in this section; thorough description will be given in a later technical report/journal version.

A. Extension to M -by- N Topologies

So far, we have focused on inferring a 2-by- N topology, which is a special but important case. M -by- N topologies can be inferred by merging one source-rooted tree topology at a time. Assume that we have inferred a k -by- N topology, $2 \leq k < M$. To add the $(k+1)^{th}$ source, we need to identify each joining point of S_{k+1} and S_i , $1 \leq i \leq k$, for each receiver, on a single logical link in the k -by- N topology (defined by all the branching points). Therefore, we need to apply Alg. 1 (or Alg. 2) to S_{k+1} and any one (in the best case) or all (in the worst case) of the current k sources. Thus, for example using Alg. 1, the number of quartets required to identify the M -by- N topology is between $(M-1)(N-1)$ and $\binom{M}{2}(N-1)$.

B. Extension to Noisy Case

So far, we have considered the noiseless scenario, where the answer to each query is the correct quartet type. One can extend the algorithms to deal with noisy queries, using the two approaches proposed in [24]. The first one is a simple solution that applies to both Alg. 1 and Alg. 2; it repeats the query multiple times and considers the majority vote as the answer to that query. The second approach is more sophisticated and fits naturally in the GBS framework.⁹ It assigns weights to each hypothesis using a probability distribution. The initial weighting is uniform, but it gets updated after each query. The update naturally boosts the probability measure of the hypotheses that agree with the answer to the query. At the end, the hypothesis with the largest weight is selected. We can adopt this approach for Alg. 2 by incorporating the probability measures in the path updates and in computing the benefits. Using this approach, Alg. 2 can handle the noisy queries more naturally than Alg. 1. The query complexity and probability of error of both approaches have been analyzed in [24].

IX. CONCLUSION

Although active topology inference is a well-studied problem, to the best of our knowledge, this paper is the first to focus on efficient merging algorithms. We propose a greedy bottom-up approach that queries only $N-1$ quartets, which is much less than $\binom{N}{2}$ possible quartets. We also formulate the problem as multiple hypothesis testing and develop an active learning algorithm based on GBS. Comparing the two proposed algorithms in simulation, we find that the simple bottom-up algorithm is near-optimal, and comparable to the

GBS baseline in terms of the number of queries (thus measurement bandwidth), while having much lower time and space complexity; therefore it is preferable for all practical purposes.

In future work, it would be interesting to fully develop the possible extensions outlined in Section VIII and also to compare our algorithms against the optimal, computed, *e.g.*, using dynamic programming (DP), which is both challenging to formulate and would have exponential complexity.

REFERENCES

- [1] M. Rabbat, M. Coates, and R. Nowak, "Multiple source internet tomography," *IEEE JSAC*, vol. 24, no. 12, pp. 2221–2234, Dec. 2006.
- [2] M. Rabbat, R. Nowak, and M. Coates, "Multiple source, multiple destination network tomography," in *Proc. of IEEE Infocom*, Mar. 2004.
- [3] M. Coates, M. Rabbat, and R. Nowak, "Merging logical topologies using end-to-end measurements," in *Proc. of ACM IMC*, Miami, FL, Oct. 2003.
- [4] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu, "Network tomography: Recent developments," *Statistical Science*, vol. 19, 2004.
- [5] A. Anandkumar, A. Hassidim, and J. Kelner, "Topology discovery of sparse random graphs with few participants," in *Proc. of Sigmetrics'11*.
- [6] A. Krishnamurthy and A. Singh, "Robust multi-source network tomography using selective probes," in *Proc. of IEEE Infocom*, FL, Mar. 2012.
- [7] S. Ratnasamy and S. McCanne, "Inference of multicast routing trees and bottleneck bandwidths using end-to-end measurements," in *Proc. of IEEE Infocom*, New York, NY, March 1999, pp. 353–360.
- [8] N. Duffield, J. Horowitz, and F. L. Presti, "Adaptive multicast topology inference," in *Proc. of IEEE Infocom*, Anchorage, AK, April 2001.
- [9] N. Duffield, J. Horowitz, F. Presti, and D. Towsley, "Multicast topology inference from measured end-to-end loss," *IEEE ToIT*, vol. 48, Jan. 2002.
- [10] M. Coates, R. Castro, M. Gadhiok, R. King, Y. Tsang, and R. Nowak, "Maximum likelihood network topology identification from edge-based unicast measurements," in *Proc. of ACM Sigmetrics*, CA, June 2002.
- [11] N. G. Duffield and F. L. Presti, "Network tomography from measured end-to-end delay covariance," *IEEE/ACM ToN*, vol. 12, no. 6, Dec. 2004.
- [12] B. Eriksson, G. Dasarthy, P. Barford, and R. Nowak, "Toward the practical use of network tomography for internet topology discovery," in *Proc. of IEEE Infocom*, San Diego, CA, March 2010.
- [13] J. Ni, H. Xie, S. Tatikonda, and Y. R. Yang, "Efficient and dynamic routing topology inference from end-to-end measurements," *IEEE/ACM Transactions on Networking*, vol. 18, no. 1, Feb. 2010.
- [14] J. Ni and S. Tatikonda, "Network tomography based on additive metrics," <http://arxiv.org/abs/0809.0158>, Aug. 2008.
- [15] M. Choi, V. Tan, A. Anandkumar, and A. Willsky, "Learning latent tree graphical models," *JMLR*, vol. 12, pp. 1771–1812, May 2011.
- [16] P. Sattari, A. Markopoulou, and C. Fragouli, "Multiple source multiple destination topology inference using network coding," in *NetCod*, 2009.
- [17] P. Sattari, C. Fragouli, and A. Markopoulou, "Active topology inference using network coding," *Elsevier PHYCOM, Special Issue on Network Coding*, March 2012, <http://dx.doi.org/10.1016/j.phycom.2012.02.006>.
- [18] C. Fragouli, A. Markopoulou, and S. Diggavi, "Topology inference using network coding," in *Proc. of Allerton Conf.*, Monticello, IL, Sept. 2006.
- [19] G. Sharma, S. Jaggi, and B. Dey, "Network tomography via network coding," in *Proc. of ITA Workshop*, San Diego, CA, Feb. 2008.
- [20] B. Huffaker, D. Plummer, D. Moore, and K. Claffy, "Topology discovery by active probing," in *Proc. of SAINT*, Nara, Japan, Jan. 2002.
- [21] B. Cheswick, H. Burch, and S. Branigan, "Mapping and visualizing the internet," in *Proc. of USENIX ATC*, San Diego, CA, June 2000.
- [22] R. Govindan and H. Tangmunarunkit, "Heuristics for internet map discovery," in *Proc. of IEEE Infocom*, Tel-Aviv, Israel, March 2000.
- [23] N. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," in *Proc. of ACM Sigcomm*, Pittsburgh, PA, Aug. 2002.
- [24] R. D. Nowak, "The geometry of generalized binary search," <http://arxiv.org/abs/0910.4397>, Aug. 2011.
- [25] S. Dasgupta, "Coarse sample complexity bounds for active learning," in *Proc. of NIPS*, Vancouver, B.C., Canada, Dec. 2005.
- [26] R. Castro and R. Nowak, "Minimax bounds for active learning," *IEEE Transactions on Information Theory*, vol. 54, no. 5, May 2008.
- [27] R. Durbin, S. R. Eddy, A. Krogh, and G. Mitchison, *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press, 1999.
- [28] J. Pearl and M. Tarsi, "Structuring causal trees," *J. of Complexity*, 1986.
- [29] L. Hyafil and R. L. Rivest, "Constructing optimal binary decision trees is np-complete," *Info. Processing Letters*, vol. 5, no. 1, pp. 15–17, 1976.
- [30] P. Sattari, "Python implementations of bottom-up and gbs algorithms," <http://odysseas.calit2.uci.edu/doku.php/public:pegah-software>, Jul. 2012.

⁹A similar solution for Alg. 1 would be to perform the deletions and contractions probabilistically.